arXiv:1510.09083v2 [cs.CV] 13 Nov 2015

# Deep Cascaded Regression for Face Alignment

Hanjiang Lai *[1,2], Shengtao Xiao[2], Zhen Cui[2], Yan Pan †[1], Chunyan Xu[2], and Shuicheng Yan[2]

[1]School of Data and Computer Science, Sun Yat-Sen University, China
[2]Department of Electronic and Computer Engineering, National University of Singapore, Singapore

## Abstract

*We propose a novel cascaded regression framework for face alignment based on a deep convolutional neural network (CNN). In most existing cascaded regression methods, the shape-indexed features are either obtained by hand-crafted visual descriptors or by leaning from the shallow models. This setting may be suboptimal for the face alignment task. To solve this problem, we propose an end-to-end CNN architecture to learn highly discriminative shape-indexed features. First, our deep architecture encodes the image into high-level feature maps in the same size of the image via three main operations: convolution, pooling and deconvolution. Then, we propose "Shape-Indexed Pooling" to extract the deep features from these high level descriptors. We refine the shape via sequential regressions by using the deep shape-indexed features, which demonstrates outstanding performance. We also propose to learn the probability mask for each landmark that can be used to choose the initialization from the shape space. Extensive evaluations conducted on several benchmark datasets demonstrate that the proposed deep framework shows significant improvement over the state-of-the-art methods.*

## 1. Introduction

Face alignment is to automatically localize the facial key points, such as eyes, mouth, nose and chin. Due to its relevance to many facial analysis tasks like facial recognition [6], expression recognition [1] and facial pose estimation [34], face alignment has attracted increasing interest.

Many methods have been proposed for face alignment [34, 18, 27, 24]. Among them, the cascaded regression method [27] has emerged as one of the leading approaches for accurate face alignment. Given an initial shape $S_0$, the algorithm refines the shape by estimating a shape increment $\Delta S$. In particular, a shape increment at stage $k$ is calculated as:

$$\Delta S^k = R^k \phi^k(S^k), \tag{1}$$

where $\phi^k(S^k)$ is the **shape-indexed features** [7] which depends on the current estimated shape $S^k$. $R^k$ is the linear regression matrix. The shape is refined as $S^{k+1} = S^k + \Delta S^k$.

Despite the success, there are still some limitations in the existing methods for cascaded regression. The main limitation is that the shape-indexed features used by most of the cascaded regression methods are fixed hand-crafted features, e.g., SIFT [16], or features learned from the shallow models. These fixed hand-crafted features may not be optimal for the specific face alignment task since they are extracted in the unsupervised fashion and are not learned during the process of face alignment. To learn the features, Cao et al. [5] proposed to jointly learn the feature mapping and the linear regression matrix by a tree-based regression. Ren et al. [20] presented a learning-based local binary feature via random forest.

Different from the above methods, in this paper, we propose an end-to-end deep convolutional neural network (CNN) for face alignment. The proposed architecture is shown in Fig. 1. We propose to use a deep convolutional network to simultaneously learn the shape-indexed features and the regressors. First, we make some structural modifications for the existing CNN architectures, including firstly reducing the size of activations through the pooling layers and then enlarging the activations through the deconvolution layers [29]. For an $H \times W$ image, the output feature maps of our network are also with the size of $H \times W$. With this, the mapping between image coordinates and CNN output coordinates is straightforward. Then, we propose a "Shape-Indexed Pooling" layer to extract the shape-indexed features from the last deconvolution layer of our network, which is similar to the "Spatial Pyramid Pooling" [9]. Due to the proposed deep shape-indexed features, our method shows significant performance gains over the methods based on fixed hand-crafted features. Moreover,

*This work was mainly performed when Hanjiang Lai was staying in NUS.

†Corresponding author: Yan Pan, email: panyan5@mail.sysu.edu.cn.

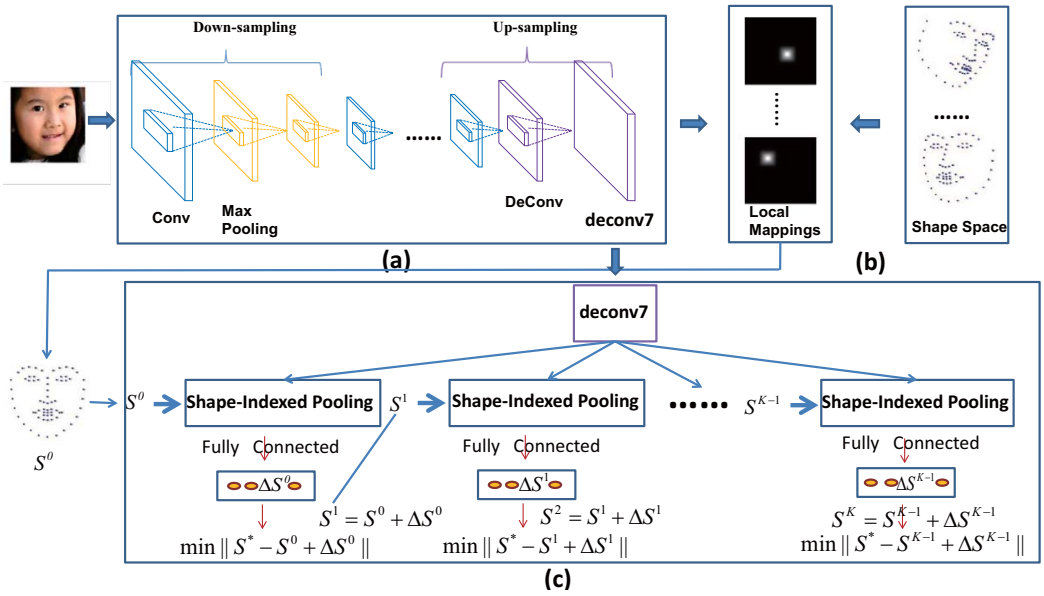


Figure 1. Overview of the proposed deep architecture for face alignment. Given an input image, our architecture (a) encodes the image into high level descriptors, e.g., deconv7, in the same size of the image, (b) learns local mapping functions to find the initial shape, and then (c) extract the shape-index features based on the deconv7 layer and the current estimated shape. The deep features are used to jointly learn sequential linear regressions for the final shape.

we propose to learn $p$ feature maps, each of which corresponds to one landmark. Each of the feature maps is a probability matrix in the same size of the input image, indicating the probability of each pixel belonging to that landmark. With these probability maps, we can find the initialization from the shape space that contains diverse shapes. With the found initialization, we refine a shape via sequential regressions, in which the fully connected layers are used to capture the linear mapping functions.

Our contributions in this paper can be summarized as follows. First, we propose a new CNN architecture for face alignment, in which down-sampling is carried out via two pooling layers and then up-sampling is performed via two deconvolution layers. Second, we design a new pooling way, referred to as "Shape-Indexed Pooling", to extract the deep features based on the current estimated shape. Finally, we show that the proposed network significantly outperforms the sate-of-the-art methods on three databases. Note that we do not use any external data sources and report the results of 68 facial landmarks.

## 2. Related Work

Many methods have been proposed for face alignment, including CLM-based methods [18], AAAM-based methods [8], Regression-based methods [27] and Deep-learning based methods [24].

**Cascaded regression methods**. The cascaded regression approaches predict the facial shape in a cascaded manner, which start from an initial shape, and refine the shape stage by stage. A representative method is the Supervised Descent Method (SDM) [27]. SDM uses SIFT features extracted around the current shape and solves a series of linear least square problems. Ren et al. [20] proposed to learn the local binary features with random forest, which achieves state-of-the-art results and much faster prediction speed. Global SDM (GSDM) [28] is an extension of SDM which divides the search space into regions of similar gradient directions.

In the above methods, the initial shape is typically provided by the mean shape. If the mean shape is the frontal face, the results may be poor for some faces with large pose variations. Cao et al. [5] proposed to use different initializations to run the algorithm and take the median of all predictions as the output. Zhu et al. [33] proposed a coarse-to-fine method which can find the initialization from the whole shape space.

**Deep learning based methods.** Deep learning has shown excellent performance in various visual applications due to its strong representation power. Luo et al. [17] proposed a hierarchical face parsing method based on the deep belief network (DBN). Sun et al. [24] proposed three-level cascaded deep convolutional networks. Each level contains several convolutional networks. The first level gives the initial estimate and the following levels refine the initial estimate and give more accurate results. Wu and Ji [26] explored deep belief networks and the 3-way restricted Boltzmann machine for the facial feature tracking. Zhang et al. [31] showed that learning the face alignment together

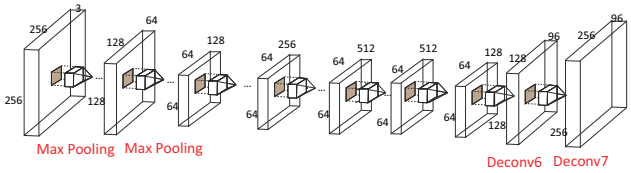


Figure 2. Deep convolution sub-network based on the variant of the VGG-19.

with some correlated tasks, e.g., head pose estimation, can improve the detection robustness.

# 3. Our Method

## 3.1. Architecture

Fig. 1 shows the detailed configuration of the proposed deep network. Our network is composed of three modules. The first module is deep convolutional sub-network. The deep network corresponds to the feature extractor which transforms the input image to the feature maps in the same size of the input image. The second module is to find the initial shape from the shape space. In this module, we design a new loss function with label distribution to learn local mapping functions. The initialization is selected using these local masks and the shape space. The third module is cascaded regression. It uses the shape provided by the second module as the initial shape. After that, we use "shape-indexed pooling" to extract the CNN features and refine the shape via sequential regressions.

## 3.2. Deep Convolution Sub-network

Convolutional neural networks have shown excellent performance in various visual problems due to their representation power. In this module, we propose a new architecture that transforms the input image to the feature maps in the same size of the input image.

Our network is based on the variant of the VGG-19 layer net [22] as shown in Fig. 2. We make the following structural modifications. The first modification is to remove the last three max pooling layers (pool3, pool4, pool5) and all the fully-connected layers (fc6, fc7, fc8). The second is to add two deconvolution layers [15, 19] (deconv6, deconv7). *Deconvolution*, which is also called *backwards convolution*, reverses the forward and backward passes of the convolution. It is used for up-sampling.

Our sub-network has 16 convolution layers, 2 max pooling layers and 2 deconvolution layers. In the convolution layers, we zero-pad the layer's input with $\lfloor k/2 \rfloor$ zeros on all sides, where $k$ is the kernel size. By this padding, the input and the output can have the same size. The pooling layer filters the input with kernel size of $2 \times 2$ and stride of 2 pixels, which makes the size of the output be half of the input. In the deconvolution layer, the input is filtered

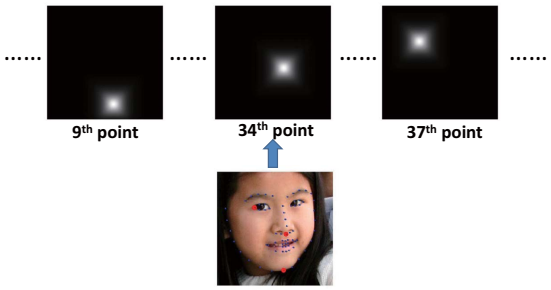


Figure 3. Learning local mappings for each landmark.

with 96 kernels of size $4 \times 4$ and a stride size of 2 pixels, which makes the size of the output be 2 times of the input. Suppose that the input image size is $H \times W$. After passing through two pooling layers, it becomes $H/4 \times W/4$, and then is upsampled to $H \times W$ via two deconvolution layers.

**Discussions.** Traditional networks, such as AlexNet [12], reduce the size of activations by repeated operations of pooling and convolution. The downsampling may cause the loss of spatial accuracy, especially for the face alignment problem. For instance, VGG19 takes a $224 \times 224$ image as the input and the output size of the $\text{Conv}_5$ feature maps is $14 \times 14$. That is, a pixel index $(x, y)$ in the $\text{Conv}_5$ feature maps has a receptive field centered on the pixel $(16x, 16y)$ in the input image. If two facial key points are very close, e.g., two points in the left eye, it is hard to discriminate them because these two points are mapped to the same pixel in the CNN feature maps.

Another CNN architecture is to remove the pooling layers and keep the size of all layers the same (e.g. $H \times W$) by repeated operations of convolution. Although it can solve the loss of the spatial accuracy problem, it has a very high time complexity due to the large image size.

On the contrary to the above two architectures, our network enlarges the activations though the deconvolution operations. The size of the image is kept in the last feature map. Hence, our network can solve the loss of the spatial accuracy problem. Also our network can reduce the size of the layers, which can greatly reduce the time complexity.

After this module, the network is divided into two branches: one is finding the initialization, and the other is the cascaded regression.

## 3.3. Initialization Searching

Many algorithms, e.g. [27], use the mean shape as the initialization. In this paper, we propose a simple method to find the initialization from the shape space instead of using a specified initialization.

We firstly construct $N$ candidate shapes $\{S^1, S^2, \cdots, S^N\}$, which cover a wide range of the shapes including different poses, expressions, etc. To obtain these candidate shapes, we simply run k-means on

the training set to find $N$ representative shapes. Then we find the initialization from these candidate shapes.

Second, we show how to find the initialization. A good initialization should be close to the ground truth shape, whereas in the testing image the latter is unknown. In this paper, we propose a simple method to give the predicted shape. Suppose the predicted shape is $S$. Then we can find the initialization as $S^0 = \arg\min_{i=1,\cdots,N} ||S^i - S||^2$.

**Learning local mapping functions.** Now the problem is how to find the predicted shape. After the sub-network, we add a new conv layer (conv8) which has $l$ channels. Since the size of conv8 is the same as the input image, a pixel index $(x, y)$ in the conv8 feature maps can be mapped to same pixel in the input image. We make the $i$-th feature map predict the location of the $i$-th facial key point as shown in Fig. 3. The location of the largest value in the $i$-th feature map is the facial key point.

We denote the $i$-th feature map as $A^i \in \mathbb{R}^{H\times W}$ and the $i$-th facial key point as $(x_i, y_i)$, where $H$ / $W$ is the height / width of the feature map. Let $\hat{A}^i_{jk} = \frac{e^{A^i_{jk}}}{Z}$, where $Z = \sum_{jk} e^{A^i_{jk}}$. $\hat{A}^i$ is the probability matrix, and $\hat{A}^i_{jk}$ indicates the probability of this pixel belonging to the landmark.

A good possibility matrix should preserve the following information: (1) the probability for the index $(x_i, y_i)$ should be the largest; and (2) the farther it is away from $(x_i, y_i)$, the smaller the probability should be. Therefore, we introduce a new ground truth probability matrix $Q^i \in \mathbb{R}^{H\times W}$, which is calculated as $Q^i_{jk} = 0.5^{\max(|x_i-j|,|y_i-k|)}$ and satisfies the above two principles. Finally, a normalization process is calculated as to ensure $\sum_{jk} Q^i_{jk} = 1$.

Since $Q^i$ and $\hat{A}^i$ are two probability matrices, we propose to use the softmax loss to quantify the dissimilarity between the predicted probability matrix $\hat{A}^i$ and the ground truth probability matrix $Q^i$, which is defined as

$$\min -\sum_{jk} Q^i_{jk} \log(\hat{A}^i_{jk}). \tag{2}$$

Note that the local mapping functions can not only help find the initialization, but also help learn the parameters of our deep network. Since the predicted shape in this subsection is obtained by only using the local information, it needs to be refined for the more accurate results.

### 3.4. Cascaded Regression

In this subsection, we present how to learn the cascaded regression based on our deep neural network. Formally, suppose we are given a training set $\{I_i, S^*_i\}_{i=1,\cdots,n}$, where $I_i$ is the $i$-th face image and $S^*_i = (x_1, y_1, \cdots, x_p, y_p)$ is the corresponding facial key points. $(x_j, y_j)$ denotes the location of the $j$-th point and $p$ is the number of facial key points.

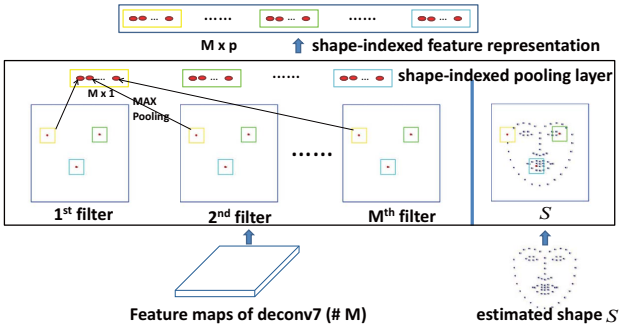


Figure 4. Shape-Indexed Pooling layer. We extract the deep shape-indexed features based on the current estimated shape S and the deconv7 layer.

Given the initial shape $S^0$ from the second module, the cascaded regression approach predicts the facial shape $S^k$ in a cascaded manner. We update the shape sequentially with the iterations $k = 0, \cdots, K-1$ as follows:

$$S^{k+1}_i = S^k_i + \Delta S^k_i, \tag{3}$$

where $S^k_i$ is the current estimate of the $i$-th image after the $k$-th iteration. $\Delta S^k_i = R^k \phi^k(S^k_i) + b^k$, which is the linear combination of the feature vector $\phi^k$ plus a biased term $b^k$. $\phi^k(\cdot)$ represents the shape-indexed features from the current predicted landmarks. $R_k$ is a linear regression matrix and $b_k$ is a biased term. A series of $\{R_k, b_k\}$ are learned in the training stage, via the following loss function:

$$\min \sum_{i=1}^{n} ||S^*_i - S^k_i - \Delta S^k_i||^2. \tag{4}$$

The problem (4) is the well-known least square problem. This step aims to regress the $S^k_i$ to a shape $S^{k+1}_i$ closer to the hand-labeled landmarks $S^*_i$.

To integrate the deep network into the cascaded regression, two problems need to be solved. The first is how to corporate the learned parameters $\{R^k, b^k\}$ into our network and the second is how to extract the features $\phi^k(S^k_i)$.

Since $\{R^k, b^k\}$ represents the linear matrix and the bias, it can be rewritten as the fully-connected layer, where $R^k$ is the weight filter and $b^k$ is the bias filter in the fully-connected layer. Hence, we add $K$ fully-connected layers in our network. The parameters $\{R^k, b^k\}$ can be learned via the back propagation algorithm.

In most existing works, $\phi^i_k$ is the hand-crafted features, e.g., HOG and SIFT. These features are fixed when performing the cascaded regressions. In this paper, we propose a new way to extract the features similarly to [9], which is referred to as "Shape-Indexed Pooling" (SIP).

The SIP layer needs two inputs. The first input is feature maps. Note that different from SIFT and HOG which take an image as input, our proposed method uses the deep

network to encode the image into high level descriptors, which increases the representational power of the image. In this paper, we choose the deconv7 as the feature layer. The second input is the current estimated shape. We extract the local features for all landmarks. Fig. 4 shows the network structure for extracting the features based on the shape $S_i^k$. More specially, for each point $(x_l, y_l) \in S_i^k$, we encode it as a bounding box where its top-left and bottom-right coordinates are $[x_i - b, y_i - b, x_i + b, y_i + b]$, with $b$ representing an integer to decide the width of the box. And we pool the responses of each filter in the local region $[x_i - b, y_i - b, x_i + b, y_i + b]$ using the max pooling. Hence, the output for each point is an $M$-dimensional vector and $M$ is the number of filters in the deconv7 layer. For all landmarks, we concatenate all the vectors into to a long $Mp$-dimensional vector. This procedure can be viewed as the deep "shape-indexed" feature.

# 4. Experiments

## 4.1. Datasets and Data Augmentation

We conduct extensive evaluations of the proposed method on three benchmark datasets.

- **LFPW** [3]: The Labeled Face Parts in-the-Wild (LFPW) database contains 1,287 images downloaded from the Internet. Due to some invalid URLs, we evaluate the performance on 811 training and 224 test images provided by [21].
- **HELEN** [13]: It contains 2,330 annotated images downloaded from the Flickr. We use 2,000 images as the training set and 330 images as testing.
- **300-W** [21]: The 300-W dataset consists of 3,148 training images from the LFPW, HELEN and the whole AFW [34]. It performs testing on three parts: common subset, challenging subset and the full set. The common subset contains 554 images from LFPW and HELEN databases and the challenging subset contains 135 images from IBUG. The full set is the union of them (689 images).

We conduct evaluations on 68 points (provided by [21]) on the LFPW, HELEN and 300-W datasets.

**Data augmentation.** We train our models only using the data from the training data without external sources. To reduce overfitting on the training data, we employ three distinct forms of data augmentation to artificially enlarge the dataset.

The first form of data augmentation is to generate image rotations. We do this by rotating the image into different angles including $\{\pm 30, \pm 25, \pm 20, \pm 15, \pm 10, \pm 5, 0\}$.

The second form of data augmentation is to disturb the bounding boxes, which can increase the robustness of our results to the bounding boxes. We randomly scale and translate the bounding box for each image.

The third form of data augmentation is mirroring.

After the data augmentation, the number of training samples is enlarged to 52 times, which is shown in Table 1.

Table 1. The number of training samples after data augmentation.

| LFPW | HELEN | 300-W |
|---|---|---|
| 42,172 | 121,160 | 163,696 |

**Evaluation.** We evaluate the alignment accuracies by two popular metrics, the mean error and the cumulative errors. The mean error is measured by the distances between the predicted landmarks and the ground truths, normalized by the inter-pupil distance, which can be calculated by

$$\text{mean error} = \frac{1}{n}\sum_{i=1}^{n}\frac{||S_i - S_i^*||^2}{pD_i} \tag{5}$$

where $S_i$ is the predicted shape and $S_i^*$ is the ground-truth shape for the $i$-th image. $D_i$ is the distance between two eyes. $p$ is the number of landmarks and $n$ is the total number of face images.

We also report the cumulative errors distribution (CED) curve, in which the mean error larger than $l$ is reported as a failure. Let $e_i = \frac{||S_i - S_i^*||^2}{pD_i}$, and CED at the error $l$ is defined as

$$\text{CED} = \frac{N_{e \le l}}{n}, \tag{6}$$

where $N_{e \le l}$ is the number of images on which the error $e_i$ is no higher than $l$.

**Implementation details.** We implement the proposed method based on the open source Caffe [10] framework, which is an efficient deep neural network implementation. We first crop the image using the bounding box with the $0.2W$ padding on all sides (top, bottom, left, right), where $W$ is the width of the bounding box. Then we resize it to the size of $256 \times 256$ pixels. The number of candicate shapes is set to $N = 5000$. We set $K = 8$ and $b = 3$. Our network is trained by stochastic gradient descent with 0.9 momentum. The min-batch size of images is 4 and the weight decay parameter is 0.0001. The network's parameters are initialized with the pre-trained VGG19 model.

## 4.2. Comparison with State-of-the-art Algorithms

The first set of experiments is to evaluate the performance of the proposed method and compare it with several state-of-the-art algorithms.

Zhu et al. [34], DRMF [2], ESR [5], RCPR [4], SDM [27], Smith et al. [23], Zhao et al. [32], GN-DPM [25], CFAN [30], ERT [11], LBF [20], cGPRT [14] and CFSS [33] are selected as the baselines. Table 2 shows the comparison results of mean error on the three datasets.

Table 2. Mean Error on three databases.

| LFPW Dataset | | Helen Dataset | | 300-W Dataset | | | |
|---|---|---|---|---|---|---|---|
| Methods | 68 pts | Methods | 68 pts | Methods | Common | Challenging | Fullset |
| Zhu et.al | 8.29 | Zhu et.al | 8.16 | Zhu et.al | 8.22 | 18.33 | 10.20 |
| DRMF | 6.57 | DRMF | 6.7 | DRMF | 6.65 | 19.79 | 9.22 |
| | | ESR | 5.70 | ESR | 5.28 | 17.00 | 7.58 |
| RCPR | 5.67 | RCPR | 5.93 | RCPR | 6.18 | 17.26 | 8.35 |
| SDM | 5.67 | SDM | 5.50 | SDM | 5.57 | 15.40 | 7.50 |
| | | | | Smith et.al | | 13.30 | |
| | | | | Zhao et.al | | | 6.31 |
| GN-DPM | 5.92 | GN-DPM | 5.69 | GN-DPM | 5.78 | | |
| CFAN | 5.44 | CFAN | 5.53 | CFAN | 5.50 | | |
| | | | | ERT | | | 6.40 |
| | | | | LBF | 4.95 | 11.98 | 6.32 |
| | | | | LBF fast | 5.38 | 15.50 | 7.37 |
| | | | | cGPRT | | | 5.71 |
| CFSS | 4.87 | CFSS | 4.63 | CFSS | 4.73 | 9.98 | 5.76 |
| CFSS Practical | 4.90 | CFSS Practical | 4.72 | CFSS Practical | 4.79 | 10.92 | 5.99 |
| OURS | **4.57** | OURS | **4.25** | OURS | **4.19** | **8.42** | **5.02** |

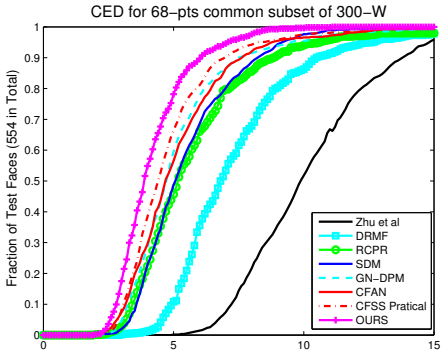


Figure 5. CED curves on common subset of 300-W.

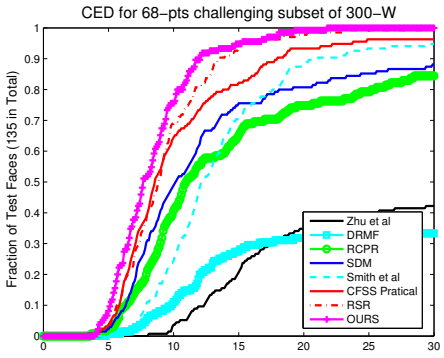


Figure 6. CED curves on challenging subset of 300-W.

It can be observed that the proposed method performs significantly better than all previous methods in all settings. Specifically, on the LFPW dataset, our method obtains a mean error of 4.57, which gives an error reduction of 0.3 compared to the second best algorithm. On HELEN, our method shows an error reduction of 0.38 in comparison with the second best method. On the 300-w dataset, the mean error of our method gives a reduction of 0.54 / 1.56 on the common set and the changeling set.

Table 3. The results of the local mapping functions.

| LFPW | Helen | 300-W | | |
|---|---|---|---|---|
| 68 pts | 68 pts | Common | Challenging | Fullset |
| 7.60 | 5.81 | 5.74 | 16.56 | 7.86 |

Fig. 5 shows the CED curves for different error levels on the 300-W dataset. Again, for all error levels, our method yields the highest accuracy and beats all the baselines. For instance, the proposed method shows a relative increase of 23% on the 300w common set compared to the second best algorithm. The example alignment results of our method are shown in Appendix.

One main reason for the good performance of our method is that instead of using traditional hand-crafted visual features (SIFT, HOG), it uses the deep network to learn the image representations and extracts the deep shape-indexed features.

### 4.3. Further Analyses

**Results of local mapping functions.** Table 3 shows the results of local mapping functions. The results are not good enough due to the fact that the mapping function only considers the local information. However, it gives us a good initial estimation compared with the mean shape.

**Effects of different input sizes and networks.** This set of experiments is to explore the effects of different deep sub-networks and different sizes of input images. We show the results of two different types of frameworks: VGG-S [1]

[1] https://github.com/BVLC/caffe/wiki/Model-Zoo

Table 4. Mean Errors on different networks and different input sizes.

| | 300-W | | |
|---|---|---|---|
| | Common | Challenging | Fullset |
| VGG-S (128x128) | 5.19 | 9.75 | 6.08 |
| VGG-S (256x256) | 4.51 | 9.64 | 5.51 |
| VGG-19 (128x128) | 4.81 | 9.22 | 5.67 |
| VGG-19 (256x256) | 4.19 | 8.42 | 5.02 |

and VGG-19, where VGG-S is a small convolutional neural network and it only has 5 convolution layers. Note that some modifications are made on the two frameworks as described above. We also report the results with two different input sizes: 256x256 and 128x128. Table 4 shows the comparison results, from which it can be seen that: (1) the larger network or larger input size will lead to better results; (2) our framework is capable of exploiting different types of characteristics, i.e., accuracy or speed, by using different sub-networks or different input sizes.

## 5. Conclusion

In this paper, we proposed an end-to-end deep-network-based cascaded regression method for face alignment. In the proposed deep architecture, an input image is firstly encoded into high level descriptors in the same size of the input image. Based on this representation, we proposed to learn a probability map for each facial key point and use these probability maps to find the initialization for the cascaded regression. Finally, we proposed a shape-indexed pooling layer to extract the deep features based on the deconv7 layer and the current estimated shape. Empirical evaluations on three datasets show that the proposed method significantly outperforms the state-of-the-arts.

## References


[1] A. B. Ashraf, S. Lucey, J. F. Cohn, T. Chen, Z. Ambadar, K. M. Prkachin, and P. E. Solomon. The painful face–pain expression recognition using active appearance models. *Image and vision computing*, 27(12):1788–1796, 2009. 1

[2] A. Asthana, S. Zafeiriou, S. Cheng, and M. Pantic. Robust discriminative response map fitting with constrained local models. In *CVPR*, pages 3444–3451, 2013. 5

[3] P. N. Belhumeur, D. W. Jacobs, D. J. Kriegman, and N. Kumar. Localizing parts of faces using a consensus of exemplars. *TPAMI*, 35(12):2930–2940, 2013. 5

[4] X. P. Burgos-Artizzu, P. Perona, and P. Dollár. Robust face landmark estimation under occlusion. In *ICCV*, pages 1513–1520, 2013. 5

[5] X. Cao, Y. Wei, F. Wen, and J. Sun. Face alignment by explicit shape regression. *IJCV*, 107(2):177–190, 2014. 1, 2, 5

[6] C. Chen, A. Dantcheva, and A. Ross. Automatic facial makeup detection with application in face recognition. In *ICB*, pages 1–8, 2013. 1

[7] P. Dollár, P. Welinder, and P. Perona. Cascaded pose regression. In *CVPR*, pages 1078–1085, 2010. 1

[8] X. Gao, Y. Su, X. Li, and D. Tao. A review of active appearance models. *IEEE T SYST MAN CY C*, 40(2):145–158, 2010. 2

[9] K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. In *ECCV*, pages 346–361. 2014. 1, 4

[10] Y. Jia. Caffe: An open source convolutional architecture for fast feature embedding. *h ttp://caffe. berkeleyvision. org*, 2013. 5

[11] V. Kazemi and J. Sullivan. One millisecond face alignment with an ensemble of regression trees. In *CVPR*, pages 1867–1874, 2014. 5

[12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012. 3

[13] V. Le, J. Brandt, Z. Lin, L. Bourdev, and T. S. Huang. Interactive facial feature localization. In *ECCV*, pages 679–692. 2012. 5

[14] D. Lee, H. Park, and C. D. Yoo. Face alignment using cascade gaussian process regression trees. In *CVPR*, pages 4204–4212, 2015. 5

[15] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. *arXiv preprint arXiv:1411.4038*, 2014. 3

[16] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004. 1

[17] P. Luo, X. Wang, and X. Tang. Hierarchical face parsing via deep learning. In *CVPR*, pages 2480–2487, 2012. 2

[18] B. Martinez, M. F. Valstar, X. Binefa, and M. Pantic. Local evidence aggregation for regression-based facial point detection. *TPAMI*, 35(5):1149–1163, 2013. 1, 2

[19] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. *arXiv preprint arXiv:1505.04366*, 2015. 3

[20] S. Ren, X. Cao, Y. Wei, and J. Sun. Face alignment at 3000 fps via regressing local binary features. In *CVPR*, pages 1685–1692, 2014. 1, 2, 5

[21] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *ICCVW*, pages 397–403, 2013. 5

[22] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3

[23] B. M. Smith, J. Brandt, Z. Lin, and L. Zhang. Nonparametric context modeling of local appearance for pose-and expression-robust facial landmark localization. In *CVPR*, pages 1741–1748, 2014. 5

[24] Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for facial point detection. In *CVPR*, pages 3476–3483, 2013. 1, 2

[25] G. Tzimiropoulos and M. Pantic. Gauss-newton deformable part models for face alignment in-the-wild. In *CVPR*, pages 1851–1858, 2014. 5

[26] Y. Wu, Z. Wang, and Q. Ji. Facial feature tracking under varying facial expressions and face poses based on restricted boltzmann machines. In *CVPR*, pages 3452–3459, 2013. 2

[27] X. Xiong and F. De la Torre. Supervised descent method and its applications to face alignment. In *CVPR*, pages 532–539, 2013. 1, 2, 3, 5

[28] X. Xiong and F. De la Torre. Global supervised descent method. In *CVPR*, pages 2664–2673, 2015. 2

[29] M. Zeiler, G. Taylor, and R. Fergus. Adaptive deconvolutional networks for mid and high level feature learning. In *ICCV*, pages 2018–2025, Nov 2011. 1

[30] J. Zhang, S. Shan, M. Kan, and X. Chen. Coarse-to-fine auto-encoder networks (cfan) for real-time face alignment. In *ECCV*, pages 1–16. 2014. 5

[31] Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Facial landmark detection by deep multi-task learning. In *ECCV*, pages 94–108. 2014. 2

[32] X. Zhao, T.-K. Kim, and W. Luo. Unified face analysis by iterative multi-output random forests. In *CVPR*, pages 1765–1772, 2014. 5

[33] S. Zhu, C. Li, C. C. Loy, and X. Tang. Face alignment by coarse-to-fine shape searching. In *CVPR*, pages 4998–5006, 2015. 2, 5

[34] X. Zhu and D. Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *CVPR*, pages 2879–2886, 2012. 1, 5